arXiv:1605.08242v1 [cs.LG] 26 May 2016

Neighborhood Sensitive Mapping for Zero-Shot Classification using Independently Learned Semantic Embeddings

Gaurav Singh[1], Fabrizio Silvestri[2], and John Shawe-Taylor[1]

[1] UCL, UK
gaurav.singh.15@ucl.ac.uk
[2] Yahoo, UK

Abstract. In a traditional setting, classifiers are trained to approximate a target function $f : X \to Y$ where at least a sample for each $y \in Y$ is presented to the training algorithm. In a zero-shot setting we have a subset of the labels $\hat{Y} \subset Y$ for which we do not observe any corresponding training instance. Still, the function f that we train must be able to correctly assign labels also on $\hat{Y}$. In practice, zero-shot problems are very important especially when the label set is large and the cost of editorially label samples for all possible values in the label set might be prohibitively high. Most recent approaches to zero-shot learning are based on finding and exploiting relationships between labels using semantic embeddings. We show in this paper that semantic embeddings, despite being very good at capturing relationships between labels, are not very good at capturing the relationships among labels in a data-dependent manner. For this reason, we propose a novel two-step process for learning a zero-shot classifier. In the first step, we learn what we call a *property embedding space* capturing the "*learnable*" features of the label set. Then, we exploit the learned properties in order to reduce the generalization error for a linear nearest neighbor-based classifier.

Keywords: Zero-Shot Classification, Textual Data, Property Embeddings, Learning Embeddings, Nearest Neighbor Classification

Introduction

One of the most prominent areas of research in machine learning is *classification*. In a traditional setting, the problem of classification consists of training a model to approximate a target function $f : X \to Y$ where at least a sample for each $y \in Y$ is presented to the training algorithm. A common problem faced by many systems, especially in their early stage, is the absence of training instances for all possible classes. In such cases, a traditional classifier cannot, in fact, assign class labels that have not been observed in the training set. This is one of the main reasons for the growth of a research area, *zero-shot classification*, studying how classification can be done also using unseen labels. In a zero-shot

setting [18], we are given a subset $\hat{Y} \subset Y$ of the labels for which we do not observe any corresponding training instance. Still, the function f that we train must be able to correctly assign labels also on $\hat{Y}$. There is a growing need for using classification systems in order to automate different online and offline tasks, like sentence tagging, image captioning, task labeling, etc. Based on the definitions given by [18], we address the following general research question: "*Given a semantic encoding of a large set of concept classes, can we build a classifier to recognize classes that were omitted from the training set?*"

In order to correctly deal with unseen classes, one possibility is to establish a relationship between seen and unseen classes. [18] propose a method for zero-shot learning using a linear model based on creating semantic embeddings for words based on co-occurrences of labels in dictionaries and based on human feedback about labels properties. The dramatic progress made in the area of semantic embeddings like word2vec ([14]) have provided a method for encoding the semantic meaning into words. Therefore, classification tasks where the label set is made up of meaningful words can be used to establish such inter-label relationships. From these relationships a plain nearest neighbors approach is taken to identify reasonable labels, also among those not seen in the training set. These methods that are aimed at "directly learning" a map to a semantic space may suffer from two major problems. First, the target semantic spaces may not be "learnable", i.e. they may not correspond to "learnable feature(s)" in the data. In other words, these semantic embeddings may correspond to contextual features of the label that may or may not be learnable from the given description of a sample. It basically means that given two instances, say x_i and x_j, they can be mapped to semantic space vectors b_i and b_j using a linear map W only if $\|x_i - x_j\|$ is proportional to $\|b_i - b_j\|$. If the semantic space vectors are learned completely independently (which is the case with semantic embeddings like *word2vec*) then there is no guarantee that the semantic vectors would be appropriate for the learning task at hand. In such cases the training error is larger and as such also the generalization error results heavily impacted. Second, nearest neighbor classifiers are not designed to take into account the neighborhood of a given label in the label space; roughly, labels that have really close neighbors should be learned with more accuracy compared to labels that are isolated in their neighborhood.

Based on the above observations, we try to address the two mentioned problems in this paper in a step-wise manner. Firstly, we need to relearn semantic embeddings based on data so that they correspond to learnable feature(s) in the data. Secondly, we need to develop a neighborhood sensitive mapping that can help reduce the risk of error in the case of nearest neighbor based classifier for zero-shot classification as used by [18].

Our main contributions in this study are the following:

- Extracting data-dependant property embeddings for labels.
- Neighborhood sensitive mapping to reduce classification error.

Related Works

There has been active research in the area of zero-shot classification in the recent past. Originally, the problem was defined in its current form by [18], where they address the problem by embedding the set of labels into a semantic space used then to extrapolate information about unseen labels. Given an object in the dataset, they firstly predict the set of semantic features (in the embedding space) corresponding to that input, and then they find the nearest class in the labels embedding. Authors also develop a generalisation bound on error for zero-shot classification using a NN-classifier. In a following study, [6] propose a multi-label max-margin classifier with applications to zero-shot. By means of a correlation matrix between different labels they are able to predict also unseen labels. The method, specifically designed for multi-label classification, explicitly reduces the hamming loss between prediction vectors associated to labels and true labels' vectors.

There are numerous practical applications where unseen images need to be classified. As a result, zero-shot classification has found considerable interest in the area of image recognition and classification. [10] develop one such approach specifically based on learning attribute for animals (like color, eating habits etc.). These attributes are not unique to a single animal and therefore can be learned from the available training data. [13] develop an approach specifically targeting images that focuses on exploiting co-occurrences of visual concepts in images for knowledge transfer. [7] propose an approach for zero-shot classification where image attributes are unreliable and use the error tendencies of the different attributes to develop a linear discriminant model. There have been many such attribute based learning methods for visual recognition that have been developed in the past: [9], [4], [19, 21, 24, 25], to name a few.

In cases where such attribute information is not easily available we need to build the semantic space for labels. Building such a space may lead to an additional overhead and it may also require an extensive knowledge of the domain in which classes are defined. For this reasons, the learned semantic space may not even be of the required high quality. A general methodology to learn word embeddings is presented by [14]. The technique, known as *word2vec*, has made a breakthrough in both efficiency and quality of the vectors learned. These vector representations were easily available for billions of words trained on terabytes of *google news* and *wikipedia* data. By using semantic embeddings learned using word2vec, [17] proposed *ConSE* a zero-shot image classification system specifically tested on image classification. ConSE uses a convolution network to embed an image into a vector space and uses a convex combination of nearest embeddings in the semantic space. It assumes that the set of predicted labels is disjoint to the set of seen labels, an assumption that is not valid in the real world and in our work. [22] proposed a linear regressor for zero-shot classification that also relies on independently learned semantic data, it is very similar to [18] in practise. [12] learns label representations from scratch without using independently generated semantic embeddings, which differentiates it from our work. [11] propose a max-margin multi-class zero-shot classifier with the assump-

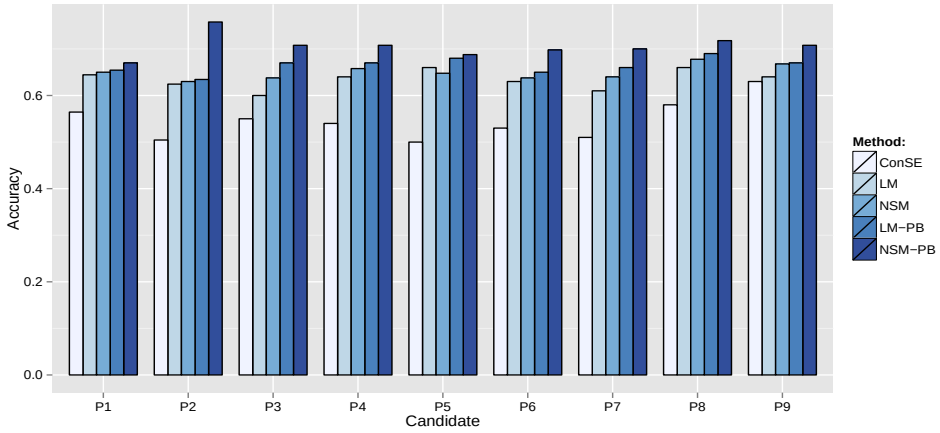

Fig. 1: Binary classification accuracy for different participants in fMRI dataset using our proposed approach (referred to as **NSM-PB**) versus LM-PB, LM, NSM and ConSE. The results are averaged over 1000 different pairs of binary classes and the results are statistically significant at p-value = 0.0016 using a two sided t-test. In this experiment, we test the ability of NSM-PB to distinguish between two novel classes as measured in [18]. The average results for NSM-PB, LM-PB, ConSE, LM and NSM: 0.7049, 0.6633, 0.5454, 0.6420 and 0.6495 respectively.

tion that unlabeled data is available during training and therefore they develop a semi-supervised approach for classification, an assumption that is not originally made in the paper defining the zero-shot classification problem and that, therefore, we are not making in this paper as well.

Proposed Method

In this section we present our two-step approach to zero-shot classification. We first describe the method to build learnable properties from the semantic representation. Then we present how to minimize the generalization error for multi-class classification using the learned properties.

In many cases data has properties that cannot be learned in a linear fashion, e.g. multi-class image classification has long involved non-linear convolution nets ([8,17]), neural nets ([3]), and non-linear SVMs ([1]). In other cases, instances are better classified using linear models. In this work, we would focus our attention on instances that contain properties which can be learned using a linear model ([5,16,20]), specifically in a zero-shot scenario.

Learning Properties

Semantic embeddings (like word2vec, but not limited to word2vec) are learned independently of the data or task of classification. In many cases, some semantic meanings encoded in these embeddings may not be learnable using the given feature(s) of the data. As an example, two animals in semantic space may be close to another word like "cute" or "faithful", now given the weight and height of an animal we can not learn a classifier that can predict such properties of the data. Again, two different varieties of cats may be very similar in the input space, but far apart in the semantic space (quite possible since semantic embeddings are independent of data). Such a situation would make it impossible to map data into semantic space without errors. Therefore, it is safe to conclude that semantic embeddings in their original form can be erroneous and unreasonable, as they are learned completely independent of the data. They may or may not be learnable given the data. Hence, we attempt to learn property embeddings from given semantic embeddings *explicitly* using data features. These property embeddings are learned based on the data while maintaining the similarities in the original semantic space (as much as possible), therefore they do not suffer from the same problem.

Previously proposed zero-shot classifiers use semantic embeddings to encode labels. In those systems the embedding is learned independently from the training samples. This is a gap we are bridging in this paper and we propose the use of what we call *property embeddings*. First of all, there is a one-to-one correspondence between labels and properties so that the i-th label is associated with the i-th property vector. In addition, property embeddings encode label vectors so that similarities among labels in the original semantic space are preserved also in the property space. Finally, objects in the training space are to be described by means of these properties. This leads us to formulate the objective J_s in Equation (1) in which the first part ensures that the property embeddings B can be learned from data X using the model W, whereas the second part ensures that the similarities that existed between different classes in the original semantic space are maintained in the property space. The model seeks to balance as much as possible the contributions of these two terms while learning property embeddings. Although, multiple instances can have the same label and therefore, we take the average of the features for all such different instances that have the same label. We denote by $\mathcal{S}$ ($\mathcal{U}$) Set of seen (unseen) labels and the subscript s (u) refers to parameters corresponding to seen (unseen) labels. Here, $X \in \mathbb{R}^{|S| \times f}$, $W \in \mathbb{R}^{f \times n^{'}}$, $B_s \in \mathbb{R}^{|S| \times n^{'}}$, $B_u \in \mathbb{R}^{|U| \times n^{'}}$, $L_u \in \mathbb{R}^{|U| \times n}$ and $L_s \in \mathbb{R}^{|S| \times n}$

$$J_s = \alpha \|XW - B_s\|^2 + (1-\alpha)\|B_s B_s^T - L_s L_s^T\|^2 + \lambda \|W\|^2 \tag{1}$$

$$W, B_s = \underset{W, B_s}{\arg\min}\, J_s \tag{2}$$

The above objective can be minimized using gradient descent. We model a slightly different objective to learn property embeddings for zero-shot (unseen) labels. Semantic embeddings provide information about relative similarity between different labels. In case of unseen labels we need to preserve the similarity

between two labels in their original semantic space because this is the only piece of information we have about such unseen labels, as we do not have any instances for such labels except the label itself. This naturally leads us to formulate the objective as:

$$J_u = \|B_s B_u^T - L_s L_u^T\|^2 \tag{3}$$

$$B_u = \arg\min_{B_u} J_u \tag{4}$$

The objective can again be optimized analytically or by using gradient descent.

Learning neighborhood-sensitive mapping to semantic space

In this section we focus our attention on developing an approach for zero-shot multi-class nearest neighbor classification using the properties learned in the previous section. Our objective is to learn a linear map $\mathcal{F}$: $X^f \longrightarrow B^{n^{'}}$ from raw input space X^f to property space $B^{n^{'}}$, such that

$$Y = \mathcal{F}(.)$$

Followed by learning a map $\mathcal{H}$: $Y^{n^{'}} \longrightarrow L$ from property space Y to the nearest label L

$$L = \mathcal{H}(Y)$$

In our case, $\mathcal{H}$ is a nearest neighbor classifier, while $\mathcal{F}$ is the linear model that maps data into property space. [18] developed a novel generalization bound for zero shot classification using a linear classifier. They used the analysis developed by [2] for nearest neighbor classifiers. They develop an upper bound (Equation (2) in [18]) on the accuracy of a linear zero-shot classifier using G_q, which is defined as follows:

$$G_q(\tau_q) = 1 - (1 - R_q(\tau_q))^n$$

where,

$$G_q(\tau_q) = P(n_q \leq \tau_q)$$

and

$$R_q(\tau_q) = P(d(q, q^{'}) < \tau_q)$$

Here, n_q is distance to nearest neighbor of predicted point q and τ_q is distance to true nearest neighbor of predicted point q. Differently from [18], we start from the upper bound to the generalization error (see Equation(2) in [18]) and instead proceed to minimize G_q as it directly contributes to the generalization error. Therefore,

$$\theta = \arg\min G_q(\tau_q|\theta) = \arg\min R_q(\tau_q|\theta) = \arg\max S_q(\tau_q|\theta)$$

$$S_q(\tau_q) = P(d(q, q^{'}) > \tau_q) = 1 - R_q(\tau_q)$$

Here, $q^{'}$ is any other class in the space.

$$S_q(\tau_q) = P(||x^{(i)}V - b_j|| > \tau_q) : i \neq j$$

V is the linear model that maps input to relearned semantic space.

$$P(||x^{(i)}V - b_j|| > \tau_q) = \frac{\sum_j [\![||x^{(i)}V - b_j|| > ||x^{(i)}V - b_i||]\!]}{|S| - 1}$$

where S is the set of all classes (seen and unseen) and $[\![P]\!]$ is the notation used to represent Iverson's brackets that evaluates to 1 iif P is true and to 0 otherwise.

Given that we normalize all property vectors i.e. $||b_i|| = 1 \ \ \forall i$

$$\left[\!\left[\left(||x^{(i)}V - b_i||^2 - ||x^{(i)}V - b_j||^2\right) > 0\right]\!\right] = \left[\!\left[\left(x^{(i)}V\left(b_i - b_j\right)^T\right) > 0\right]\!\right]$$

and we can then write

$$S_q(\tau_q) = \frac{\sum_j \left[\!\left[\left(x^{(i)}V\left(b_i - b_j\right)^T\right) > 0\right]\!\right]}{|S| - 1} \tag{5}$$

Generalizing the probability for all samples assuming independence:

$$P = \prod_q S_q \tag{6}$$

Substituting S_q from Equation (5) in Equation (6) and replacing product of constants by C

$$P = \frac{1}{C} \prod_{(k,i) \in D} \sum_j \left[\!\left[\left(x_k^{(i)}V\left(b_i - b_j\right)^T\right) > 0\right]\!\right]$$

Here, (k, i) is any instance k which has label i and D is the data. Therefore, if the k_{th} data sample has label i then it constitutes a valid pair (k, i). This formulation is required because multiple instances can have the same label.

$$\log P = \sum_{(k,i) \in D} \log \left(\sum_j \left[\!\left[\left(x_k^{(i)}V\left(b_i - b_j\right)^T\right) > 0\right]\!\right] \right) - \log C$$

Approximating $\log P$ by a cancave upper bound U:

$$\log P \leq \sum_{(k,i) \in D} \sum_j \log \left(\left[\!\left[\left(x_k^{(i)}V\left(b_i - b_j\right)^T\right) > 0\right]\!\right] \right) - \log C = U \tag{7}$$

We need to approximate the $[\![\cdot]\!]$ function with a continuous function ($\sigma(x)$) that quickly goes to 1 for $x > 0$ and approaches 0 for $x < 0$. A natural choice for approximating $[\![\cdot]\!]$ is sigmoid function, $\sigma(x) = 1/(1 + e^{-x})$.

$$U = \sum_{(k,i) \in D} \sum_j \log \left(\sigma \left(x_k^{(i)}V\left(b_i - b_j\right)^T \right) \right) - \log C \tag{8}$$

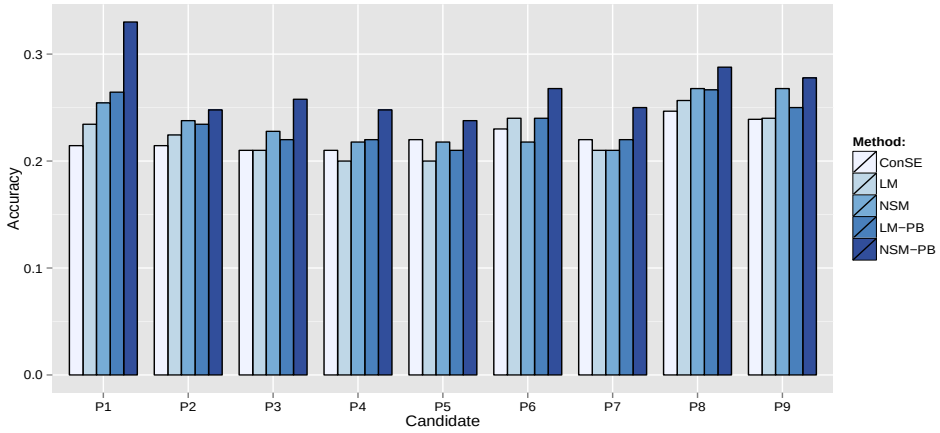

Fig. 2: The figure plots multi-class classification accuracy for different participants in fMRI dataset. The figure compares our proposed approach (referred to as **NSM-PB**) against LM-PB, LM, NSM and ConSE. The results are averaged over 1000 different set of five classes and the results are statistically significant at p-value = 0.0015 using a two sided t-test. In this experiment, we test the ability of NSM-PB to distinguish between five novel classes that have not been seen in the train set. The average results for NSM-PB, LM-PB, ConSE, LM and NSM: 0.2671, 0.2356, 0.2281, 0.22393 and 0.2349 respectively.

In addition to the above objective, we add l_2 regularization on the V and we seek:

$$V = \arg\max_{V} \sum_{(k,i)\in D} \sum_{j} \log\left(\sigma\left(x_k^{(i)} V\left(b_i - b_j\right)^T\right)\right) - \log C \\ -\lambda||V||^2 \tag{9}$$

Notice, that the objective learns a given instance based on the neighborhood of target label. It specifically enforces that labels that are embedded in popular neighborhood are learned with higher precision. For this reason, we refer to it as *Neighborhood Sensitive Mapping*.

Experimental Evaluation

In this section, we first describe the different datasets that we used in the study. Thereafter, we pose different experimental question and analyse the results in the light of the questions. In the end, we discuss the results obtained on two larger tagging datasets.

Dataset

We used three datasets, namely, fMRI dataset from [15] (as used in [18]), wiki10+ dataset as described in [26] and delicious dataset as described in [23] for experiments. We performed initial analyses on fMRI dataset and then computed results on wiki10+ dataset.

- The fMRI dataset is composed of the neural activity observed from nine human participants while looking at 60 different concrete words. These 60 words are divided into 12 different categories, like animals: bear, dog, cat, cow, horse and vehicles: truck, car, train, airplane, bicycle. Each participant was shown a word and a small line drawing of the concrete object the word represents. The participants were asked to think about the properties of these objects for several seconds while scans of their brain activity were recorded. Each sample measures the neural activity at roughly 20,000 locations in the brain. Six fMRI scans were taken for each word. We also used the same time-averaging described in [15, 18] to create a single average brain activity pattern for each of the 60 words, for each participant.
- The wiki10+ dataset from [26] contains text of *wikipedia* articles and the tags assigned by users on *delicious.com* for *url* of those articles. We used the most popular tag for an article as the label of that article. There were 20762 instances in the dataset with 5303 distinct labels. We cleaned the data of all html tags and computed tf-idf representations of the data. Afterwards, we used truncatedSVD to reduce noise and dimensionality of the data.
- The delicious dataset [23] contains features of web pages from all over the Internet with tags generated by users on those web pages as the labels. The dataset contains 500 features for each instance and 983 unique labels. It has more than 16000 instances and the features are binary, with 1 indicating the presence of the feature and 0 indicating the absence of the feature.

The semantic embeddings used in this work are 300 dimensional trained using word2vec ([14]) on *google news* dataset. These pre-trained semantic embeddings are available online for a very large number of different words.

Experiments

In these experiments, we refer to our method as NSM, that stands for *Neighborhood Sensitive Mapping*, [18] is referred to as *LM* and [17] is referred to as *ConSE*. In order to get insights into the effect of properties, we also test both NSM and LM with property embeddings. Therefore, *NSM-PB* and *LM-PB* refer to NSM using property embeddings and LM using property embeddings respectively. On the other hand, *NSM* and *LM* use the original semantic embeddings. We evaluated the proposed approach on three different research questions and compared our results to the methods in [18] and [17].

1. How well can the model differentiate between two novel classes, where neither class appears in the train dataset?

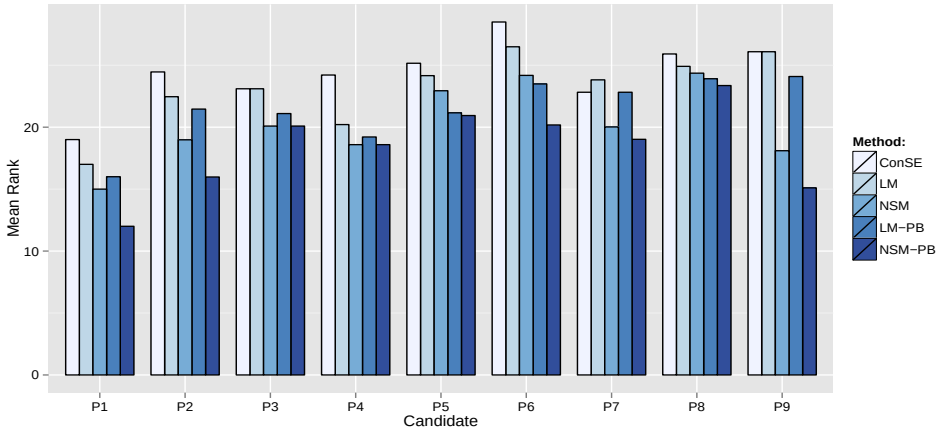

Fig. 3: The figure plots mean rank for a true novel class (not seen in the train set) in the ordered prediction list of classes for different participants. In the figure our proposed approach (referred to as **NSM-PB**) is compared against LM-PB, LM, NSM and ConSE. The difference of NSM-PB against LM and ConSE is statistically significant at p-value < 0.001 using a two sided t-test. In this experiment, we test the ability of NSM-PB to predict novel classes that have not been seen in the train set with low rank in the list of ordered prediction scores. The average results for NSM-PB, LM-PB, ConSE, LM and NSM: 18.58, 21.47, 24.36, 23.137 and 20.25 respectively.

For this task, we randomly select a set 1000 pairs, such that each pair consists of two classes. We remove these classes completely from the train dataset. Thereafter, we compute the average binary classification accuracy on all these pairs in the test dataset.

We can see that NSM-PB outperforms LM-PB, LM (Figure 1) and ConSE. The difference between NSM-PB and others is statistically significant using a t-test at p-value < 0.001. It can be seen that relearning of semantic embedding using the data, as well as modified neighborhood sensitive approach manages to better differentiate between novel classes. In fact, during our experimentation we observed that it works particularly well in case of classes that are very close in original semantic space, e.g. *hammer* and *chisel*. In such cases, the use of a direct linear mapping leads to results that are no better than random. We obtained an accuracy of 50% for *hammer* and *chisel* using LM, whereas 67% accuracy using NSM-PB.

2. How well can the model classify accurately in a multi-class classification setting, where all the test classes are absent in the train dataset?

We randomly select a set of 1000 groups such that each group consists of five different classes. We compute the average of the results of multi-class classification accuracy on all these groups. We restrict the predicted class to the classes

in the group, i.e. we try to predict the correct class from among the five classes in each group. The results are presented in Figure 2 and are statistically significant using a two sided t-test at p-value = 0.0015. It is interesting to note that even in a multi-class setting property based models outperform the semantic embedding based models, also NSM-PB outperforms ConSE, although ConSE is competitive.

We can see in Figure 2 that NSM-PB performs better than both LM and ConSE. It shows that the proposed model is better at discriminating between multiple zero-shot classes in a multi-class classification setting. We can see in Figure (3) that the mean rank of correct label in the prediction list is much lower in case of NSM-PB as compared to both LM-PB, LM and ConSE. It shows the effectiveness of NSM-PB in predicting correct labels higher in the prediction list.

3. How well can the model predict accurately in a multi-class classification setting, where the classifier has to choose from all possible classes?

For this task, we select a random class and remove it from the train dataset. Thereafter, we try to predict that class during testing from the set of all classes, including both novel and seen classes. It means that the classifier has to choose a class from among the 60 different classes present in the dataset. This is the hardest task and also closely resembles the real world situation where we do not know in advance if the instance belongs to a seen or unseen class. The results are presented in Figure 4 and are statistically significant using a t-test at p-value < 0.001. NSM-PB outperforms the baselines in this task as well (see Figure 4). We make the classifier choose from the complete set of 60 classes instead of restricting the possible set of classes. Note that we outperform LM-PB, LM, NSM and ConSE by a significant margin in terms of classification accuracy (See Figure 4). These results are as expected given that we minimize generalization error for nearest neighbor classifier. This clearly shows in the results and both LM and ConSE fall short in comparison to NSM for classification accuracy. These results clearly give us the insight regarding the usefulness of property embeddings, as LM-PB performs much better than LM and NSM-PB performs much better than NSM. In addition to property embeddings, the neighborhood sensitive mapping further improves the accuracy of classification.

K	NSM-PB	LM	ConSE
5	0.0685	0.0261	0.0283
10	0.1344	0.0392	0.0472
50	0.3590	0.0521	0.0825

Table 1: The value of accuracy for different methods on wiki10+ dataset. A given prediction is considered accurate if top-K labels in the prediction list contain the correct label. Results are statistically significant using t-test at p-value < 0.001.

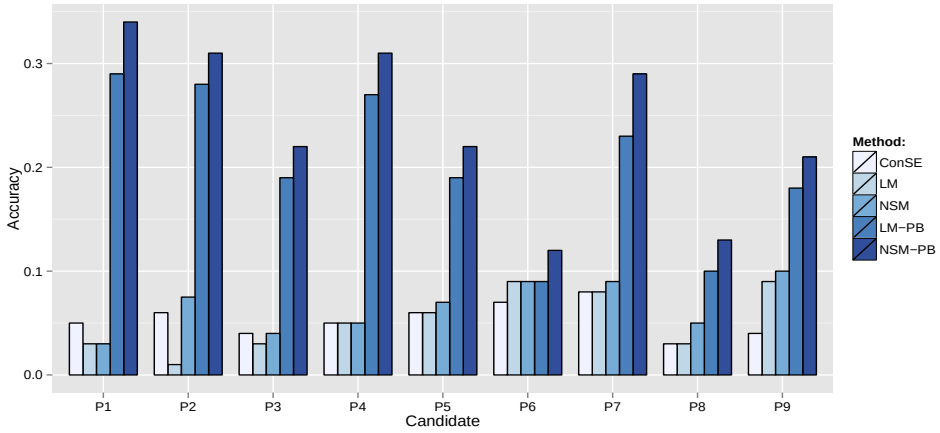

Fig. 4: The figure plots mean accuracy for a true novel class (not seen in the train set) for different participants in fMRI dataset. In the figure our proposed approach (referred to as **NSM-PB**) is compared against LM-PB, LM, NSM and ConSE. The difference of NSM-PB against LM and ConSE is statistically significant at p-value < 0.001 using a two sided t-test. In this experiment, we test the ability of NSM-PB to predict a novel class from among the set of all possible classes. The average results for NSM-PB, LM-PB, ConSE, LM and NSM: 0.2389, 0.2000, 0.0533, 0.0522 and 0.0661 respectively.

In the end, we tested NSM-PB against LM and ConSE on 2 much bigger datasets called **wiki10+** ([26]) and another tagging dataset called **delicious** [23]. The wiki10+ dataset contains 20K+ instances with more than 5000 unique labels. We created 100 different test-train splits with each test set consisting of 100 zero-shot labels, while the train set consisting of all the other labels. The classifier had to choose the correct label from the set of all possible labels, that is both train and test labels. It means that the classifier didn't have prior knowledge whether a test instance was a zero-shot label or not, which is close to the real life situation for a classifier. Since a given url can genuinely have multiple correct labels, we decided to test the accuracy of the method using the top-K selection of predicted labels. It means that the prediction was considered accurate if the top-K predicted labels contained the correct label.

The delicious dataset contains textual data of web pages along with their tags. It contains 500 features and 983 unique labels. The cardinality of the dataset is 19.020. If a given sample had more than one label then we replaced it with one row in the data with the row corresponding to one of the labels (randomly chosen). For example, if a sample x had labels (a, b) then we randomly included either (x, a) or (x, b) in the dataset. We created 100 different test-train splits with each test set consisting of 100 zero-shot labels, while the train set consisting of all the other labels. As a result the total number of labels in the dataset was

reduced to 400+. The classifier was made to choose the correct label from the set of all possible labels, as in the previous case. We decided to test the accuracy of the method using the top-K selection of predicted labels. It means that the prediction was considered accurate if the top-K predicted labels contained the correct label.

We can see in Table 1 that NSM-PB outperforms the other approaches by a significant margin for varying levels of K. We can see that NSM-PB performed respectably even for K=5, which is a very small value considering that the classifier has more than 5000 labels to choose from. We can see very similar results in Table 2 that NSM-PB again outperforms all other approaches for varying values of K. We observe that NSM-PB performs better for smaller values of K, as well as larger values of K, which is impressive given that the total number of classes to choose from was quite large. The performances were far better than for a random classifier.

K	NSM-PB	LM	ConSE
5	0.0509	0.0320	0.0433
10	0.0924	0.0492	0.0822
20	0.1486	0.0721	0.1350

Table 2: The value of accuracy for different methods on delicious dataset. A given prediction is considered accurate if top-K labels in the prediction list contain the correct label. Results are statistically significant using t-test at p-value < 0.001.

Reproducibility

We share the codes used for the given experiments at `http://bit.ly/1RCNlwR`. The values of α and λ in Equation (1) that give best results are 0.1 and 0.5 respectively. In case of ConSE, we use a multi-class logistic regression classifier for predicting class probabilities. The values of parameter T (i.e. number of top-T nearest embeddings for a given instance) in ConSE that gave best result was 5. The dimensionality of the learned property embeddings in the experiments was 10.

Conclusion

In many web applications, a number of instances are labelled constituting a subset of the entire set of possible labels. Most of these systems utilize semantic embeddings to learn correlation between different labels. Recently, there have been some advancements in the area of semantic embeddings leading to wide popularity and easy availability of pre-trained semantic embeddings. In this paper, we use these pre-trained semantic embeddings to improve a linear zero-shot

classifier in a multi-class classification setting. We first illustrate the problems with linear zero-shot classification systems that use semantic embeddings. We also highlight problems of classifiers that are not sensitive to neighborhood of a label in semantic space. Afterwards, we develop an insight into extracting property embeddings that better correspond to learnable features in the data. We show that NSM and LM both work better when supplied with such property vectors instead of semantic vectors. In addition, we show that NSM, which is based on minimizing an approximation to generalization error performs better than LM. Similarly, NSM-PB outperforms LM-PB, ConSE and LM, which proves the effectiveness of using property embeddings and neighborhood sensitive approach to zero-shot classification.

Future Work

In this work, we have observed that relearning semantic embeddings can really improve the quality of classification results. Although, In this paper we restricted the analysis to extracting linear property embeddings. In the future, we will extend the work in the direction of learning property embeddings from non-linear data with max-margin non-linear classification. We will also create a multiple kernel version of the proposed approach. We believe that it would drastically improve the quality of the results for complex data with multiple non-linearities.

References

1. Olivier Chapelle, Patrick Haffner, and Vladimir N Vapnik. Support vector machines for histogram-based image classification. *Neural Networks, IEEE Transactions on*, 10(5):1055–1064, 1999.
2. Paolo Ciaccia and Marco Patella. Pac nearest neighbor queries: Approximate and controlled search in high-dimensional and metric spaces. In *Data Engineering, 2000. Proceedings. 16th International Conference on*, pages 244–255. IEEE, 2000.
3. Dan Ciresan, Ueli Meier, and Jürgen Schmidhuber. Multi-column deep neural networks for image classification. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 3642–3649. IEEE, 2012.
4. Ali Farhadi, Ian Endres, Derek Hoiem, and David Forsyth. Describing objects by their attributes. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 1778–1785. IEEE, 2009.
5. David Guillamet, Jordi Vitria, and Bernt Schiele. Introducing a weighted non-negative matrix factorization for image classification. *Pattern Recognition Letters*, 24(14):2447–2454, 2003.
6. Bharath Hariharan, SVN Vishwanathan, and Manik Varma. Efficient max-margin multi-label classification with applications to zero-shot learning. *Machine learning*, 88(1-2):127–155, 2012.
7. Dinesh Jayaraman and Kristen Grauman. Zero-shot recognition with unreliable attributes. In *Advances in Neural Information Processing Systems*, pages 3464–3472, 2014.

8. Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012.
9. Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Learning to detect unseen object classes by between-class attribute transfer. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 951–958. IEEE, 2009.
10. Christoph H Lampert, Hannes Nickisch, and Stefan Harmeling. Attribute-based classification for zero-shot visual object categorization. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 36(3):453–465, 2014.
11. Xin Li and Yuhong Guo. Max-margin zero-shot learning for multi-class classification. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, pages 626–634, 2015.
12. Xin Li, Yuhong Guo, and Dale Schuurmans. Semi-supervised zero-shot classification with label representation learning. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.
13. Thomas Mensink, Efstratios Gavves, and Cees GM Snoek. Costa: Co-occurrence statistics for zero-shot classification. In *Computer Vision and Pattern Recognition (CVPR), 2014 IEEE Conference on*, pages 2441–2448. IEEE, 2014.
14. Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. *Advances in Neural Information Processing Systems 26*, 2013.
15. Tom M Mitchell, Svetlana V Shinkareva, Andrew Carlson, Kai-Min Chang, Vicente L Malave, Robert A Mason, and Marcel Adam Just. Predicting human brain activity associated with the meanings of nouns. *science*, 320(5880):1191–1195, 2008.
16. Andriy Mnih and Ruslan Salakhutdinov. Probabilistic matrix factorization. In *Advances in neural information processing systems*, pages 1257–1264, 2007.
17. Mohammad Norouzi, Tomas Mikolov, Samy Bengio, Yoram Singer, Jonathon Shlens, Andrea Frome, Greg S Corrado, and Jeffrey Dean. Zero-shot learning by convex combination of semantic embeddings. *International Conference on Learning Representations*, 2013.
18. Mark Palatucci, Dean Pomerleau, Geoffrey E Hinton, and Tom M Mitchell. Zero-shot learning with semantic output codes. In *Advances in neural information processing systems*, pages 1410–1418, 2009.
19. Devi Parikh and Kristen Grauman. Relative attributes. In *Computer Vision (ICCV), 2011 IEEE International Conference on*, pages 503–510. IEEE, 2011.
20. Calyampudi Radhakrishna Rao and Helge Toutenburg. *Linear models*. Springer, 1995.
21. Marcus Rohrbach, Michael Stark, and Bernt Schiele. Evaluating knowledge transfer and zero-shot learning in a large-scale setting. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1641–1648. IEEE, 2011.
22. Bernardino Romera-Paredes and PHS Torr. An embarrassingly simple approach to zero-shot learning. In *Proceedings of The 32nd International Conference on Machine Learning*, pages 2152–2161, 2015.
23. Grigorios Tsoumakas, Ioannis Katakis, and Ioannis Vlahavas. Effective and efficient multilabel classification in domains with large number of labels. In *Proc. ECML/PKDD 2008 Workshop on Mining Multidimensional Data (MMD08)*, pages 30–44, 2008.

24. Naman Turakhia and Devi Parikh. Attribute dominance: What pops out? In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 1225–1232. IEEE, 2013.
25. Felix X Yu, Liangliang Cao, Rogério Schmidt Feris, John R Smith, and Shih-Fu Chang. Designing category-level attributes for discriminative visual recognition. In *Computer Vision and Pattern Recognition (CVPR), 2013 IEEE Conference on*, pages 771–778. IEEE, 2013.
26. Arkaitz Zubiaga. Enhancing navigation on wikipedia with social tags. *arXiv preprint arXiv:1202.5469*, 2012.